

World Class Production and Inventory Management

Darryl V. Landvater

THE OLIVER WIGHT COMPANIES

Oliver Wight®

Oliver Wight Publications, Inc.
5 Oliver Wight Drive
Essex Junction, Vermont 05452-9985

his time. It was Jim who first motivated me to become involved in MRP II.

Finally, my hope is that this book makes a contribution which is in addition to the profits that companies achieve through the use of more competitive tools. I hope it does something to elevate the quality of life for those thousands of people who today struggle to manage manufacturing companies without the proper tools to do their jobs.

Darryl Landvalter
October 1992
Williston, Vermont

Contents

CHAPTER 1	3
The Emergence of a Profession	
CHAPTER 2	19
Major Themes in PIM	
CHAPTER 3	35
Rethinking Business as Usual	
CHAPTER 4	51
Sales and Operations Planning	
CHAPTER 5	71
Demand Management	
CHAPTER 6	97
Master Production Scheduling	
CHAPTER 7	123
Material Requirements Planning (MRP)	
CHAPTER 8	159
Capacity Planning	
CHAPTER 9	179
Plant Scheduling	
CHAPTER 10	203
Supplier Scheduling	
CHAPTER 11	217
Implementation: The Proven Path	
EPilogue	257
The Road Less Traveled	261
INDEX	

us a quarterly commitment at the start of each model year. The dealers tell us the number of units they're going to take and give us a forecast of desired units by quarter. They don't have to take those commitments, but if they do, we give them a cash rebate. We've tried to develop a program that helps us identify the needs of our dealers, and consequently our future production needs."

Alternatively, when a company cannot influence customer demand, it can at least manage it. Identifying marketplace demand patterns, making well-thought-out projections, and taking a realistic look at the market may not alter demand, but these activities allow a company that respond more effectively to customer demand than a company that doesn't do as effective a job in each of these areas. The classic example is a financially managed company that sorts off with a target number for earnings per share. It then uses this number to calculate a target for profits, and finally translates this back into the sales numbers. This entire process takes place in a vacuum, without a focus on the marketplace.

Of course, this is not to say that financial goals aren't important—they are. But so is the marketplace. One of the overwhelming objectives for any company is to meet its customer service objectives. A customer-driven company will see demand management as a tool for satisfying its customers—a way to provide higher levels of customer service and increased customer responsiveness.

John Lewis, plant manager at Continental Can Co., found that his company's work on demand management provided a set of questions to ask the salespeople and the customers. "For example," explains Lewis, "You said you were going to sell 500,000 units and you only sold 410,000. Are the other 90,000 units going into the next month, or are they dropping out completely? With the data, the salespeople can get more comfortable sitting down with the customers. They say, 'You said you were going to take 500,000 of this, and you only took 410,000. This is the pattern. You're always overestimating what you'll need. Is there a reason for this?' Many times the customer will respond, 'I didn't know whether you were really going to be able to make them.' I was buying from XYZ Company, and the only way I could get 400,000 from XYZ was to say 'I needed 600,000.'"

WHY IS DEMAND MANAGEMENT IMPORTANT?

If you can stabilize the top-level plans (production plan and master production schedule) through demand management, the effects will be felt throughout the entire organization. Well-managed companies have found that stabilizing the top-level plans creates enormous productivity improvements: all functions operate more smoothly and effectively. For example, when customer orders are promised on an existing schedule that meets the customer's needs, salespeople make better use of their time than they do when they must constantly act as go-betweens in negotiating deliveries. Manufacturing people focus on hitting the plans rather than frantically trying to deal with the latest changes to the schedule. The plant can better use equipment and labor, and supplier delivery performance improves because suppliers aren't chasing moving targets. Finally, quality improves because the company is not trying to push most of the month's shipments out the door at the last minute.

Demand management is not a panacea. No one has a crystal ball that can perfectly anticipate the "zigs and zags" of the marketplace—there will always be surprises. But demand management can help you better serve your customers and make life easier for everyone in your company.

With this as background, let's look at the different subjects in demand management:

- sales planning/forecasting
- order promising
- distribution center demands (or branch warehouse demands)
- inplant demands
- service or spare parts demands

SALES PLANNING/FORECASTING

The difference between sales planning and forecasting is like the difference between "Sic 'em" and "C'mere." With sales planning, the "Sic 'em" questions are: "What is our action plan to make the desired sales

happen?" "What are our advertising and direct mail programs?" "Who are the customers we'll call on?" "What's our pricing?" In short, it's a proactive stance. For most people, forecasting is a passive exercise. People take a guess at the future and hope for the best.


A proactive sales plan is more likely to happen, and if it doesn't, an analysis quickly shows what worked and what didn't. Perhaps the sales promotion was more successful than planned, but the dealer program was less effective. With this information, salespeople can formulate a new set of plans and go back to the marketplace.

Tom Socio, supervisor of planning and scheduling at Krautramer-Brinson, explains the kind of progress his company has made. "Our forecast accuracy has gone from 62 percent to 84 percent. In the feedback from field salespeople, it's evident that our customer service has increased, in terms of both our ability to turn product around and our ability to respond to upsets in the business. It's enabled us to get back market share that we once lost, and we're actually getting market share in some new areas. It's mostly just discipline, pulling in sales, project management, creating the environment to have a lot of open communication, formalized meetings, formalized schedules for turning in forecasts."

"We were a typical operation, getting forecasts on an irregular basis. The production manager would review the forecasts and scale them down because he didn't believe them. In all likelihood, the sales group had increased the forecasts, knowing he would reduce them. There was no forward vision—everything was extremely short term. We were also caught in a fiscal-year-planning trap. Early in the fiscal year, things were fine, but as we went into the ninth, tenth, and eleventh months we still had no forecast for the first months of the next fiscal year. Now we've expanded that horizon to a rolling twelve-month forecast."

The approach that Tom referred to is shown in Figure 5.1. While the approach is somewhat different in every company, there are six common elements:

Figure 5.1 Sales Planning Process



1. *Document the assumptions.* The process of sales planning starts with documenting assumptions. Unfortunately, many organizations fail to separate the assumptions and the numbers. Typical assumptions are the answers to such questions as: "What do we think the economy is going to do?" "What's happening with housing starts, oil prices, interest rates, exchange rates, and the like?" They can be very specific, such as "What do we think our competitors will do in terms of new product introduction, pricing, advertising, sales activity, and so on?" They can also be internal: "What are we going to do in terms of pricing, advertising, and so on?"

2. *Convert assumptions into numbers.* The assumptions are typically documented by experienced people in the company, generally top-level managers and sales management. Someone, typically called the demand manager or demand planner, then takes these assumptions and turns them into numbers—i.e., the sales plan. The demand planning position is new to many organizations. Typically, the demand manager or planner is located in the marketing department.

Once a company's approach changes from "C'mere" to "Sic 'em," there's a lot of work to be done in all areas of demand management. The demand planner's responsibilities involve sales planning as well as order promising, distribution, interplant demands, and service parts demand.

3. *Reconcile numbers to create a sales plan by product family.* In the area of sales planning, the demand manager or planner is responsible for working with both the sales and marketing groups to help them develop their proactive plans to make the sales plan happen as well as to reconcile their numbers. The sales organization is looking at the sales plan and how to make it happen: "What activities are needed?" "What are the potential pitfalls and how would we avoid them?" Marketing is looking at total demand, market shares, product life cycles, and the like. This larger market perspective needs to be reconciled with the sales plans generated by the sales and management groups.

As you can see, the sales plan is both a top-down and bottom-up document; a number of different groups get involved with putting the plan together. Each has a different perspective and makes a valuable contribution. And in the end, the process of consensus has to be reached for the sales plan to be a company plan.

4. *Break down the sales plan by item.* Once those questions are answered, the sales plan, still at the family level, must be broken down by product. Sometimes this is done with historical information, such as "In the past, this product has had 5 percent of family sales." For example, 12.5 horsepower lawn tractors account for 25 percent of all lawn tractor sales. But again, some assumptions come into play—"Why would it be the same? Why would it be different from last year? Has anything changed in the market that would mean this percentage ought to be different?"

At this point, the sales plan can be calculated for individual items. One way this can be done is to take the sales plan for the family of items and multiply it by the expected percentage for each item in the family. If the sales plan anticipates sales of 1,600 lawn tractors and the 12.5 horsepower tractor accounts for 25 percent of sales, then the forecast for this model should be 400.

A number of companies rely on software that forecasts an item's future sales from its sales history. Forecast developed this way need to be verified in two ways. First, the future may not be like the past (it

rarely is)—do these differences mean the forecast for this item should be changed? Also, how does the historical forecast compare with the sales plan for the family or item? If there is a significant disagreement, the differences need to be reconciled. It may be that the historical forecast is reconciled to the total forecast for the family of products because the future is expected to be different from the past. On the other hand, a historical forecast that adds up to more or less than the sales plan may open some eyes and cause the sales plan to be revised.

5. *Measure actual sales against the sales plan.* "We just started measuring the forecast," explains Ed Wohlwend, manager of inventory control/central planning at Hillshire Farm & Kahn's. "Everybody had an opinion. Fred Miller of Industrial Engineering has a saying: 'Without data you're just another person with an opinion.' Once we started looking at the data, some things were obvious just from looking at the graphs."

Since a forecast is never "right," meaningful measurement of a sales plan requires a tolerance. Let's say a sales plan calls for sales of 1,000 units per month and actual sales last month were 1,090. This might represent a statistical fluke, and no action is needed on this sales plan. However, if the actual sales were 1,500, then something is clearly happening. It could be a statistical peak, in which case we'd expect a corresponding valley; and over time the sales would average around 1,000 per month. On the other hand, this difference could mean something is happening in the marketplace, and the sales and marketing people should zero in on this and find out what has changed. A tolerance is needed to focus attention on those differences that should be reviewed and to distinguish them from minor deviations.

Forecasting has few if any absolutes. The best way to look at forecasting is to consider it a process in which you take the difference between what you predicted would happen and what actually happened, then make it smaller. Forecasting spare parts demands for large generators, for example, is mostly a matter of mathematics. Every so many hours the bearings need to be replaced. Based on the number of machines in the fleet, the age, and usage, a forecast for parts can be calculated with a reasonable degree of accuracy. On a monthly or quarterly basis, forecast accuracy of 80 or 90 percent might be possible.

High-fashion items, on the other hand, are anyone's guess. The first guess is whether or not an item will catch on in the marketplace. If it does, the next guess is to what degree, and, finally, when will the

popularity fade? Both the rise and fall can be dramatic, and the period in between is like walking a tightrope. Sales forecasts for these items are far less accurate than those for the bearings mentioned above.

Consequently, the primary measurement of sales planning is an improvement in the accuracy with which the sales plan predicted actual sales. Did the difference between the plan and what actually happened get smaller? For the turbine bearings, did the accuracy rise from 80 percent to 88 percent? For fashion earrings, did the accuracy change from 55 to 60 percent? If so, then some new goals should be set.

Most companies measure sales plan accuracy by family and by individual item. Generally, the accuracy of the forecast for the family of items is better than for the individual items. It's easier to forecast the number of tires that people will buy than the number of specific types of tires (snow tires, all-season tires, tires of different widths, and so on). While it's important to measure both, in the end a manufacturing company makes specific items, not families. If the sales plan for the family is fairly accurate, but the mix is off significantly (for example, you've sold 200 percent of plan for product #2 and 50 percent of plan for product #3), then it's likely that you've bought the wrong materials and may have the wrong capacity.

6. *Review assumptions.* A difference between actual sales and the sales plan may indicate a change in the marketplace. It may be necessary to reforecast the item or family. The first step in this process is to reexamine the assumptions that were used to develop the sales plan. Is the economy recovering faster than expected? Are housing starts increasing? Was our promotion more successful than we predicted? The assumptions may need to be revised, and new sales plan numbers calculated.

The marketing group needs to look at these situations from an overall perspective. Is the product going through its life cycle faster than anticipated? Is the product being replaced by another technology? Is its market share slipping?

The sales group should also look at any sales plan that is out of tolerance—its members are in a good position to answer questions like, "What's happening with our major customers?" "In competitive situations, are we experiencing more losses than expected?" "Are we doing better than anticipated?"

In any of these instances, the most fundamental issue is to understand

the situation and revise the numbers accordingly. In short, it's a problem-solving activity that seeks to find the root cause by updating the assumptions. Consensus may not be easy to reach, but each of the group plays an important role, and each brings key information to the process.

In many companies, an impediment to this is the tendency to "shoot" the forecaster. Whenever anything goes wrong, someone can always point to the forecast and say, "If only we had an accurate forecast."

In some cases, manufacturing people have unrealistic expectations about forecast accuracy. Sales people, having been shot at before, develop their own ways to present a smaller target. One sales manager's flippant advice to a new salesman was "Son, give me a date or a number, but never both." In other situations, the forecast can be the flash point for conflict between manufacturing and sales. Each is convinced that the other group is the problem. This leads to a losing situation and impedes the rounds of finger pointing that generate more heat than light. A better approach is to create realistic expectations and work together as a team to make improved forecasts.

A fundamental principle of forecasting is that the further out into the future you forecast, the less accurate your forecast will be. This only makes sense, because more things can change the further out you go in time. Yet many companies have long lead times for material and capacity, and this often means significant commitments are made for material and capacity based on their forecasts. When they decide to implement Just-in-Time, their total lead times drop significantly. This in turn means that instead of making commitments on what they think is going to happen in week 40, they make commitments for what they think is going to happen in week 20, or even week 2. The forecast for week 40 and 50 is going to be more accurate than the forecast for week 40 and 50 more accurate commitments are made. These help to reduce excess material and capacity, and also take some of the pressure off the forecasters to be accurate so far into the future. So, unlikely as it may seem at first, Just-in-Time is a significant help in the sales planning process.

ORDER PROMISING

The second major subject in demand management is order promising. Many companies promise customer orders based on standard lead time—for example, a piece of capital equipment such as a shear is

promised for delivery to new customers in six weeks, bottles of vitamin C are said to be available "off the shelf," and so on. The problem is, a desired lead time of 6 weeks may be pushed out to 10 weeks by a flood of orders.

Likewise, if you say that bottles of vitamin C are available off the shelf and you just used up all the stock, vitamin C is no longer off the shelf. It's just gone from being an off-the-shelf product to one like the shear that will be promised for some date in the future. Customers have shown repeatedly that they can ramp up product faster than manufacturers can make it. Alternatively, if the ordered lead time for the shear is six weeks, sales have been slow, and the product can be delivered in four weeks, why not take advantage of the situation?

Good customer service means selling the customer the truth, which in turn requires good information. While you will not always be in a position to be able to give customers what they want, you should (1) at least give customers dates that they can count on, even if they aren't the dates the customer wants, and (2) consider changing the master production schedule if you're violating the customer service objectives you've agreed on. (Refer to Chapter 6 for an explanation of master production schedule changes.)

One crucial tool for making effective order promises is the available-to-promise (ATP) calculation. Available-to-promise is a very simple calculation that uses the master production schedule, the on-hand inventory, and existing orders. Available-to-promise shows how much of a particular product is available to promise to customers, and when it can be safely promised. This seems remarkably simple, and you'd expect most companies to operate this way, but that's not the case.

Figure 5.2 is an example of an available-to-promise calculation showing:

1. Customer orders that have already been promised. In this case, 100 of this product have been promised to customers in week 2, 30 in week 3, and another 20 in week 4.
2. On-hand inventory, in this case 10 units of this product.
3. Scheduled production of this product by date. In this case, 200 are scheduled to be completed in week 2 and another 200 in week 4.

Figure 5.2 Available-to-Promise Calculation

On Hand = 10	Time Period				
	1	2	3	4	5
	Customer Orders				
			100	30	20
	Scheduled Production				
			200	200	
ATP			10	80	260/260

The ATP calculation shows what's available by week for new customer orders. If a customer called today, it would be possible to ship 10 units right away. These are the 10 units on hand. The production of 200 units scheduled to be completed in week 2 will be used to satisfy the customer orders for 100 in week 2 and 30 in week 3 (leaving 70 of the 200 available to promise to new customers). The production of 200 scheduled to be completed in week 4 will be used to satisfy the customer orders for 20 in week 4 (leaving 180 of the 200 available to promise to new customers).

The calculation shown in Figure 5.2 is called a *cumulative available-to-promise calculation*. This is because the 10 on hand are added to the 70 available in week 2 to give a total of 80 available-to-promise units in weeks 2 and 3. Likewise, the 80 available in week 3 are added to the 180 available in week 4 to give a total of 260 available-to-promise units in weeks 4 and 5.

While ATP arithmetic is simple, it is nonetheless quite powerful. ATP answers the question "When can I promise a customer order?" If a customer calls and says, "I need 10 as soon as possible. When can they be shipped?" the answer is "Right away." If a customer calls and says, "I need 80 as soon as possible. When can they be shipped?" the answer is, "Next week." And if a customer calls and says, "I need 260 as soon as possible. When can they be shipped?" the answer is, "Week 4. But we can give you 10 right away, 80 next week, and the remainder in week 4." Today, with technologies that provide accurate on-hand balances and on-line ATP calculations, sales, marketing, scheduling, or production

people can easily perform order promising. Salespeople generally do the order promising if the customer expects a delivery date immediately. If the nature of the business is such that some engineering work is required before a promise date can be given to the customer, then order promising might be done by another group.

Some companies have additional features in their order promising systems. For example, they will take an entire order and look at the earliest date when they can ship all of the items in the order. They may be able to ship the first item next week, the second in two weeks, and the third in three weeks. Depending on the type of business and product, the company might choose to break the order into separate shipments or ship everything in week 3. The same logic is used with a product with many options, say, a car. If a customer wants a car with a V-6 engine, an AM/FM/CD player, and an interior trim package, the order promising logic identifies the earliest date when all three options are available and promises the delivery at that time.

In many organizations, available-to-promise is an on-line function, so that as customer orders are taken, the ATP is immediately updated and the system gives a clear, up-to-date picture of when delivery can be promised. If the ATP for a particular item is outside the customer service objectives (the next order cannot be promised until week 8, and six weeks is the objective), the order entry people should notify the master scheduler and request a change to the master production schedule. The objective of the change would be to make the master production schedule consistent with the customer service objectives. This might mean increasing the master production schedule, or moving scheduled production up to an earlier date.

If the master production schedule cannot be changed, then it's necessary to influence demand. In the long term, a company wants to make what it can sell, but there are times when, in the short term, it may be necessary to sell what can be made. There are a number of ways to influence demand in these situations, such as running a promotion on the products for which the company can handle the demand.

Setting these customer service objectives isn't always as simple as it sounds. A common problem is setting objectives for high levels of customer service, but not being willing to do what it takes to achieve them. If a product is to be delivered off the shelf, then it's necessary to invest in inventory. If the delivery objective is 24 hours, then it may be necessary to improve your manufacturing process so you can make the

product quickly, and/or to invest in some inventory. It's necessary to define customer service objectives in terms of lead time to the customer, which will determine finished goods and work-in-process inventory levels using the current manufacturing processes. Either accept these levels, change the manufacturing processes, or revise the customer service objectives.

Improving the manufacturing process is what Just-in-Time is all about. Companies that used to promise customer orders eight weeks in the future can, using Just-in-Time, manufacture and deliver a product in a few days. Companies that were stocking hundreds of different products have eliminated their finished goods inventory and can now manufacture to the customer order in 24 hours. This is an integration issue that needs to be examined in a company contemplating a Just-in-Time implementation.

Most companies measure order promising by tracking how well they do in meeting their customer service targets. This is different from the kinds of external measurements that all companies should be doing. These external measurements look at how well the company is doing in meeting customers' expectations. For example, a company's customer service objective may be delivery in four to eight weeks. An internal measurement would determine how many orders were actually delivered in more than eight weeks or less than four. The results may indicate the need to revise the production plan or master production schedule. The results may also indicate that a significant number of orders are being expeditied for special delivery (which generates significant chaos). The external measurement would detail how many customers expect delivery in less than four weeks. The results may indicate a need to revise the customer service objectives. For example, it may be necessary to revise the customer service objectives to delivery in three to five weeks.

DISTRIBUTION CENTER/BRANCH WAREHOUSE DEMANDS

"We used to sit and wait for the inevitable wave at the end of each business year," explains Bob Wagner, distribution manager for Digital Business Corporations' Westminster distribution center. "We didn't know how big it was, but we knew it was coming and couldn't take any chances. So we'd bring in an army of temporary

it paid off, but other times the wave wasn't as big as we expected, and a lot of the money was wasted."

"We had to do the same thing with the trucking companies," adds Dick McGee, Business Group manager of ORU/CSMO (Operations Business Unit/Computer Systems Manufacturing Distribution) at Digital. "We couldn't tell these vendors exactly how far the trucks we'd need or when we'd need them, so we'd just start with the maximum. Of course, when we went to negotiate contracts with them we were at a serious disadvantage—they knew that whatever we told them would be wrong."

All this was before DEC adopted distribution resource planning (DRP). "The search for a system like DRP was born out of a crisis," says Paul Mantos, DEC's senior manufacturing consultant. "We didn't make our earnings one quarter because we missed the revenues by a day—the informal system couldn't keep up. So we had our own market crash—DEC's stock slipped 5.5 points. This was not only an embarrassment to management and a scare to employees, many of whom own stock in the company, but it was like the heart attack that saved the patient. We knew we had to do something, and do it soon."

Two years later, DEC realized the following benefits as a result of DRP during a period when order growth averaged 15–20 percent per year:

1. Inventory turns improved by 61 percent.
2. Overtime was reduced by 90 percent.
3. Revenue shipments were 33 percent higher than in the previous year.
4. Freight was reduced by 16 percent.
5. Operating expenses were reduced by \$8.5 million.

Distribution resource planning also works at Mass Merchandisers Inc., which does not have its own manufacturing operations, but purchases products and then distributes them.

Ronnie Williams, senior buyer of health and beauty aids at Mass Merchandisers, recalls what things were like before DRP: "We were basing purchase orders on five-week averages. Seasonality can bite you real quick. For example, when we get a spot of cold weather, we know that we will need a huge amount of cold remedy products from the warehouse. Using a five-week sales average, you don't pick up on those situations. You quickly get behind. If you don't get a cold remedy, you don't pick up on the slowdown. So, with a five-week average, you are always faced with very quickly bringing out of stock or very quickly being overstocked. DRP levels that out."

Steve Nelson, director of inventory management systems at Mass Merchandisers, credits DRP with helping the company to achieve significant improvements in customer service. "Prior to DRP, we would drop to the 85–88 percent level and it would take six weeks to recover. Today we have solid information that allows us to react before a problem becomes critical."

And Bob Dickson, executive vice-president and vice-president for purchasing at Mass Merchandisers, relates what happened in "retailers": "We used to get pressure because inventory levels were too high, so we would work on reducing them, but there would be a big dip in the service level. Now we have had new business brought on board along with many marketplace changes, but inventory levels have remained stable. Today our turns are consistent. Although we have added some 10,000 slow-moving SKUs (up 30 percent) over the last four years, our overall turns have improved."

Like DEC and Mass Merchandisers before DRP, many companies plan by using averages for distribution center (DC) demands (also called branch warehouse demands). Say, for example, that on average a distribution network uses 3,000 cases per week of notebook paper. On average, that's correct. But what happens if the New York, Los Angeles, and Chicago DCs all need to be resupplied this week? The shipping quantity to each of these locations is several weeks' supply. So 9,000 cases of notebook paper are needed this week. Suddenly, the average is off by a factor of three. In general, peaks and valleys create customer service problems. You may have a week with demand as high as 6,000 or

86 WORLD CLASS PRODUCTION AND INVENTORY MANAGEMENT

7,000 cases, and other weeks with demands of 0, depending on how frequently you supply the DCs and what quantities are shipped.

DRP is planning and scheduling system for distribution networks. By using a computer, a planner can accurately predict when distribution centers need to be resupplied, and accurately show the peaks and valleys in the demand pattern from the distribution network. If next week there is a demand of 7,000 cases, you can see that demand 3, 10, or even 20, is in the future and plan for it. If a distribution planner can have the weeks into the future and plan for it, the planner can have the weeks into the future and plan for it. If a distribution planner can have the weeks into the future and plan for it, the planner can have the weeks into the future and plan for it.

For example, instead of shipping 2,000 cases to the DC in Houston and 2,000 cases to the DC in Chicago this week, then being unable to meet the demand for 5,000 from Los Angeles next week, a planner could choose to increase the master production schedule to cover the peak demands. Or, the planner could cut the shipping quantities to all the DC in Houston, 1,000 to the DC in Chicago, and 2,000 to the DC in Los Angeles, and then resupply them in a week or two when more inventory is available.

The point is, there's no reason to subject your company to surprises in demand. With DRP, the demands of the distribution network are visible and can be managed.

Distribution resource planning is the logic of MRP (material requirements planning)—see Chapter 7) applied to a distribution network. This logic is generalized network scheduling logic, and it works equally well for a distribution network. Figure 5.3 shows an example of DRP for a distribution center.

The forecast is what is expected to be sold from this distribution center during a time period (in this case each week). The forecast for this type of notebook paper is 800 cases per week, except in weeks 4 through 6, when it increases to 900 per week.

In Transit are the quantities of product in transit from the central supply facility to this distribution center. The quantity is shown for the date when arrival at the distribution center is expected. In this case, 2,000 cases are expected to arrive in week 1.

Projected On Hand is an inventory projection out into the future. There are currently 1,200 cases in inventory, 800 cases are expected to be sold this week, and 2,000 cases are expected to arrive, giving 2,

Figure 5.3 DRP Example, Chicago, Notebook Paper

On Hand = 1200 Cases
Safety Stock = 1000
Transit Time = 2 Weeks
Shipping Quantity = 2000

	Weeks							
	1	2	3	4	5	6	7	8
Forecast	800	800	800	900	900	900	800	800
In Transit								
Projected	2,400	1,600	2,800	1,900	1,000	2,100	1,300	2,500
On Hand	1,200							
Planned Shipments					2,000	2,000		

projected on hand balance of 2,400 cases at the end of the week. This drops to 1,600 in week 2, and will go to 800 in week 3 if nothing is done. But 800 is below the safety stock level of 1,000 cases, so the logic in the DRP software plans a shipment to arrive in week 3, uses the specified shipping quantity of 2,000 cases, and uses the specified in-transit time of two weeks.

The result appears on the Planned Shipment line. The plan is to ship 2,000 cases from the central supply facility to the Chicago distribution center in week 1, to arrive in week 3. Similarly, shipments are planned in week 4 to arrive in week 6, and in week 6 to arrive in week 8. These planned shipments are the key to providing visibility into the distribution network.

The planned shipments for each of the different distribution centers in the network can be summarized as in Figure 5.4.

Notice how the combined demands of the distribution centers can have significant peaks and valleys. The people who do the planning for the central supply facility will find this information invaluable. In week 4, for example, a number of the big distribution centers require shipments. If the planners can see the peak demand for 8,500 ahead of time, they are in a better position to have the inventory available. If they cannot see this coming, they probably will not have enough inventory to cover a peak demand of this size.

Figure 5.4 Total Distribution Demands (Planned Shipments)

Notebook Paper

	1	2	3	4	5	6	7	8
Chicago	2,000			2,000		2,000		
New York		2,000		2,000		2,000		
Los Angeles	3,000			3,000				3,000
Houston		1,500		1,500			1,500	
Toronto				1,500		1,500		
Total	5,000	3,500	1,500	3,500	0	3,500	3,500	3,000

In a nutshell, DRP takes the forecast by item for each distribution center, the inventory, and the item data (shipping quantity, shipping schedule, transportation lead time) and calculates the planned shipments from the central supply facility to the distribution centers. These planned shipments are then accumulated and reported as demands to the master scheduler (see Chapter 6 for an example of how the distribution demands are shown on the master schedule report). The central supply facility could be a manufacturing plant, several plants, or a purchasing operation in the case of a pure distributor (like Mass Merchandisers) that does no manufacturing.

Plugging DRP into the master scheduling system involves more than just creating a technical interface. DRP makes the distribution network transparent in the sense that people can now see into the network and manage it globally. Excess inventory in one location can be transferred to another. Peak demands from the distribution network are visible at the plant weeks in advance, instead of being surprises, as before.

All this means that distribution and manufacturing people used to develop a new working relationship. Andre Martin, a pioneer in the field of DRP, says, "In the past, distribution was responsible for distributing

inventory, and the unitized assumption was that inventory was there for the taking. In reality, the inventory wasn't always there, and distribution and manufacturing were put into adversarial roles." Distribution people would say, "If manufacturing could just supply the inventory, there wouldn't be stockouts." Manufacturing people's response, on the other hand, was "If distribution wouldn't keep surprising manufacturing with peak demands, we could run the plant efficiently." The information provided by DRP allows these groups to see the constraints each has to deal with, and to work effectively as a team.

The planned shipments from the central supply facility to the distribution centers can also be used for transportation planning. The distribution planning system is capable of generating a transportation planning report, which shows the weight and cube of the items to be shipped from a central supply facility to each DC each day or week. This enables a transportation planner to take best advantage of freight rates. For example, in Figure 5.5, 3.5 railcars need to be shipped to the Chicago DC this week and 1.5 railcars need to be shipped next week. The planner might consider moving half a railcar up from next week to this week and shipping early, or taking half a railcar and moving it out to ship later.


The planner is also responsible for efficiently using the railcar or truck. If a railcar is reaching its weight limit but hasn't used its cube, the planner might choose to change the mix. If the railcar for next week uses the cube but not the weight, one way to solve the problem would be for

Figure 5.5 Transportation Planning Report

		From: Central Supply				
		To: Chicago DC				
Week		1	2	3	4	5
Weight		1,380,108	550,126	705,010	420,106	750,206
Cube		34,086	13,080	18,065	8,015	7,593

1 Railcar = 400,000 Pounds
10,000 Cubic Feet

Figure 5.6 Interplant Demand



The planning system for the customer plant provides a schedule of what's needed and when it must arrive. For purchased items, this is a purchasing schedule. For items sourced from another plant, this schedule is represented as interplant demands. These demands are sent to the supplier plant, and are shown as demand on the master production schedule report (see Chapter 6 for an example of how interplant demands are represented). In effect, these are customer orders from the customer plant. This schedule of demand from the customer plant takes the place of the forecast—an application of MRP pioneer Joe Orlicky's idea, "Don't forecast what you can calculate." In terms of interplant demand, that means there's no reason for a supplying plant to be forecasting when it can simply plug into the planning system of the customer plant and accurately predict what's needed.

Once the interplant demands have been communicated from the customer plant to the supplier plant, the work begins. The work in this case is the cooperation between the two plants. On the one hand, the customer plant may need the supplier plant to change schedules more quickly than the supplier would like. The supplier plant, on the other hand, may need to stabilize its schedules because of material capacity limitations, and the ever-present demands to run the plant efficiently.

the planner to pull some of the lighter products into the current week and use the heavier ones next week.

For a complete explanation of DRP, see *Distribution Resource Planning: Distribution Management's Most Powerful Tool* by Andre Martin (Oliver Wight Publications).

INTERPLANT DEMANDS

Interplant demands (i.e., demands for a plant's product from another plant in the company) are a continuing source of frustration in many manufacturing companies. The supplying plant is usually convinced that the "customer" plant makes unrealistic demands and does not allow enough time. The customer plant, on the other hand, typically considers the supplying plant to be rigid and inflexible.

For the most part, all of this is unnecessary. The planning system at the customer plant can be plugged into the planning system at the supplying plant, and the most up-to-date picture of demand can be communicated rapidly and efficiently between the two locations. If something changes in the demand picture is different, then the two plants can work it out by determining whether they can make the change happen or by deciding what changes, if any, need to be made.

This doesn't mean that customers are changing at a rate faster than the two plants can keep up with, but that's life in manufacturing. When customers do change their demands, however, two plants with connected planning systems can get accurate information quickly. With this information, they can then decide whether a change can or cannot be realistically made.

Formerly, the three plants of the General Dynamics Land Systems Division, manufacturer of the M1A1 tank, did not have a common master production schedule. Each plant would maximize its own schedule and build up its own inventory, resulting in production that was at one time or thirty tanks behind schedule. Today, the integration and visibility provided by MRP II have allowed the different plants to tie into a unified master production schedule for tanks, so that all three plants are working from a common game plan.

Managing interplant demand as the General Dynamics Land Systems Division has done involves connecting two planning and scheduling systems that flows out of one goes into the other, as illustrated in Figure 5.6.

Consequently, it's a process of cooperative negotiation—give-and-take, knowing what's realistic, knowing what needs to be done, and knowing how to make the plan happen.

In most cases, the two plants working together can devise a plan that meets both their needs. Figure 5.7 illustrates this process.

Let's say that the customer plant makes a change in its master production schedule to meet changing customer demand. This "explodes" down through the customer plant's planning system and creates a series of changed demands. These changes are transmitted to the supplier plant, and the planners there see these changes in demand.

It may not be necessary for the supplier plant to change its master production schedule in response to these changed demands. The supplier plant may have inventory that could be used to meet the demands. For example, it may have 500 bearing housings on hand, enough to meet an increase in demand of 100. Or it might be planning to manufacture a quantity greater than the immediate need, and this might allow the new demands to be covered without any changes to the master production schedule.

Perhaps neither is the case, and the supplying plant needs to make a

change to its schedule. Let's further assume some limitations in capacity, so the full amount of the change cannot be made. It may, however, be possible to make a partial change—perhaps half the change (50), but to accommodate (the supplying plant can provide an additional 50) half the customer plant will have to find a way to do without the other half. The company may have to reduce some order quantities, compress some lead times, or, in the end, provide only a partial delivery to the customer. In any case, there needs to be a cooperative effort on the part of both the customer and the supplying plant.

Credibility can be a significant issue in making this process work. If the supplying plant says the customer plant never really knows what it needs and when it needs it, or if the customer plant says it can never get the supplier to deliver on time so it must always ask for more than it really needs, then the process breaks down. So it's the responsibility of the planners in each of the plants to maintain accurate information and credible systems, and to work honestly with each other.


Unfortunately, in some organizations there's a disincentive to work together. It may be to one plant's advantage to "stick" another plant with inventory. ("The end of the year is coming up, we're going to have a physical inventory, so let them keep it on their books.") Or let's say that the supplier plant supplies outside customers in addition to other divisions of the company. In such a case, it's usually more profitable to ship to the outside customers and let the customer plant suffer. In the end this hurts the total profitability of the organization. If these kinds of negative incentives are in place, it's important to dismantle them so people can work together as members of the same team.

SERVICE PARTS DEMANDS

All equipment manufacturers have service or spare parts demands that need to be managed well. Not only are service parts typically an extremely profitable part of business, but companies get a reputation for either having or not having service parts available. If their reputation in this regard is poor, they can have a difficult time selling their products. Finally, if service parts are not managed well, it may be necessary to cannibalize equipment being assembled for shipment. This not only affects shipments, profits, and new customer promises, but is extremely inefficient.

Service part demands can be difficult to manage because in many

Figure 5.7 Changes in Interplant Demand



companies it's fairly typical to have unrealistic customer service objectives for service parts. For example, one capital equipment manufacturer's policy was to deliver service parts off the shelf or within two weeks—yet its customers could order service parts for a piece of equipment that hadn't been made in ten years! These items were not in inventory, and in many cases the manufacturer would have to buy the raw material and fabricate the parts. In such a situation, it was impossible to provide these parts in two weeks.

A solution is to establish categories for service parts. For example, the first category might be items that are typically ordered as service parts for equipment built in the last five years; these items would be available out of stock. Another category might be items that are not typically required to service equipment, or parts for equipment more than five years old. These items would be promised based on the lead time necessary to make the parts, including raw material purchase. If a part happened to be in stock, then it would be promised right away; if the material needed to make the part was available, then the part would be promised at the lead time necessary to fabricate it; and so on. Finally, a third category might be service parts for equipment older than ten years. In this case, some decisions would need to be made on the cost of supplying the parts and the possible advantages of upgrading or even completely replacing the equipment. Both manufacturing and sales would get involved in these situations.

Once these categories have been established, forecasts are needed only for the first category—those service parts that will be supplied from stock or raw materials to be used to make service parts. The remaining categories do not require any forecasting, since they will be manufactured to order.

The forecasts for service parts are then entered directly into material requirements planning as requirements for the items. Material planning (see Chapter 7) takes into account the total demands for the item—the demands for finished products plus the demands for service.

If all the demands cannot be met, the planner has the best picture of the situation and can either make or recommend a decision. Many times, the decision is whether to use a part on equipment that has been promised to a customer or ship the part to another customer for service. As always, the objective is to do everything possible to meet both demands. When that is not realistic, however, the decision will have to be made based on the company's customer service policy.

These situations will become visible when the available-to-promise calculation is used for service parts. The same types of ATP information can be displayed for service parts as for finished products, and the process described above can be used.

SUMMARY/CHECKUP

1. The days of accepting marketplace demand are over. That passive approach has been discarded in favor of a more active stance called demand management. Companies can control a number of influencing factors, and when they control these demand, they can at least plan for it. People are making the transition from "Can't do it" to "We can do it." Where is your company on a scale of 1 to 10 with 1 being "Can't do it" and 10 being "We can do it"? Where should you be?
2. Once a company makes the transition from the passive approach to the more active process of demand management, someone needs to fulfill the role of demand manager or demand planner. Do you have a demand manager or demand planner responsible for managing the demand stream feeding your planning and control system?
3. The process of sales planning is somewhat different for each company, but the fundamentals are the same for all:
 - a. Document the assumptions.
 - b. Convert assumptions into numbers.
 - c. Reconcile numbers to create a sales plan by product family.
 - d. Break down the sales plan by item.
 - e. Measure actual sales against the sales plan.
 - f. Review assumptions where actuals are out of tolerance.

How well does your company do each of these activities?

4. Companies no longer promise customer orders based on a standard lead time or an off-the-shelf assumption. Order promising should be done using an available-to-promise (ATP) calculation. Is such a calculation used in your company?

5. Distribution resource planning (DRP) should be used to plan the demands from distribution centers where such centers exist. The

accumulated demand from the distribution network, is shown on the master production schedule report. If you have distribution centers in your company, is DRP being used in this way?

6. Where products are shipped between plants, the planning system for the customer plant should be feeding its demands to the planning system at the supplier plant. These interplant demands should appear on the master production schedule report. If you have interplant shipments in your company, is the demand being handled in this way?

7. Equipment manufacturers have to plan for service parts demands. These should be forecast only for ship-from-stock items or raw materials used to make service parts, and should be promised using an available-to-promise calculation.

Each element of demand management is important, although not all elements apply to every company. The result of demand management is a valid and well-managed demand stream. This demand stream feeds the remainder of the planning and control system, specifically the master production schedule, which is the subject of the next chapter.

SPECIAL ACKNOWLEDGMENTS

Richard Ling and George Reimagine for their contributions to the field of demand management, and Andre Martin for his original thinking on DRP.

Chapter Six Master Production Scheduling

"The toughest thing is *not* to master-schedule," explains Frank Wenzel, director of materials and systems at Sealed Air, a manufacturer of piston rings and other combustion chamber components, "but to have management *respect* the master schedule. In the past, as in most companies, if our management got a call from a high-level officer at a customer's firm, the first reaction was to violate the master schedule and do whatever needed to be done to satisfy that customer, rather than try to manage the master schedule. That's changed. We still deliver 98 percent on time, almost across the board. But we don't see upper management whiplashing the factories anymore."

"I can think of one incident in which the owner of one of our customer companies called our president about delivery. It was a delivery that would have seriously violated our master schedule. Our company management was off site, and I was with the president. He was hot after he hung up the phone. As he was calling the plant to make sure it expedited the order, I told him, 'If you do this, you become the division expeditor. Is that what you want?' Rather than expediting it, he put the problem in the hands of the plant management and the materials group. And instead of telling them to just do it, he asked them what they *could* do. In the past, it would have been 'You'll make it happen—I don't care if you're over